Combating Uncertainty with
Fusion

A Multidisciplinary Workshop: April 22-24, 2002
Woods Hole, Massachusetts

Final Report prepared by
Diana E. Jennings, Ph.D. and Misha Pavel, Ph.D.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)		
08-09-2003	FINAL	Apr. 2002 Through Sept 2003		
4. TITLE AND SUBTITLE Combating Uncertainty with Fusion: a multidisciplinary workshop		**5a. CONTRACT NUMBER** N00014-02-1-0461		
		5b. GRANT NUMBER N00014-02-1-0461		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Jennings, Diana E Pavel, Misha		**5d. PROJECT NUMBER**		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Marine Biological Laboratory 7 MBL St Woods Hole MA 02543		**8. PERFORMING ORGANIZATION REPORT NUMBER**		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) ONR - Office of Naval Research Ballston Centre Tower One 800 North Quincy Street Arlington VA 22217-5660		**10. SPONSOR/MONITOR'S ACRONYM(S)** ONR		
		11. SPONSORING/MONITORING AGENCY REPORT NUMBER		

12. DISTRIBUTION AVAILABILITY STATEMENT

Approved for public release; unlimited distribution

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This report is a summary of a NASA/ONR-sponsored workshop, **Combating Uncertainty with Fusion** that was organized in Woods Hole in April 2002. The main purpose of the workshop was to address a class of difficult computational problems that are characterized by combining large amounts of data or datasets from diverse sources that are related in complex, stochastic, and poorly understood ways. The intent was to determine whether understanding of biological fusion processes could provide guidance to the development of robust algorithms that would alleviate the difficulties encountered in a variety of application areas including the Earth Observation System.

15. SUBJECT TERMS

Sensor fusion, information fusion, biological information processing, computational modeling, sensory integration, Bayesian

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Diana E Jennings, PhD
a. REPORT	b. ABSTRACT	c. THIS PAGE	SAR	33	19b. TELEPONE NUMBER (Include area code) 508-289-7535

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI-Std Z39-18
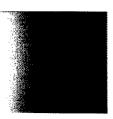
Combating Uncertainty with
Fusion

A Multidisciplinary Workshop: April 22-24, 2002
Woods Hole, Massachusetts

Final Report prepared by
Diana E. Jennings, Ph.D. and Misha Pavel, Ph.D.

Combating Uncertainty with
Fusion

Sponsored by: the Office of Naval Research, the Center for Advanced Studies in the Space Life Sciences at the Marine Biological Laboratory, Woods Hole, MA, and
NASA Ames Research Center
Meeting Chair: Misha Pavel, Ph.D.
Oregon Graduate Institute of Science & Technology

This workshop addressed a class of difficult computational problems that are characterized by the fusion of large amounts of data or datasets from diverse sources that are related in complex, stochastic, and possibly poorly understood ways. The desired output is to reduce uncertainty in various ways, ranging from supervised classification and identification of known "objects" to exploratory data analysis and pattern discovery at different time scales.

Even the most successful current approaches lack the robustness exhibited by biological systems. The fact that fusion problems are solved by biological systems, for example in pattern recognition or during tasks requiring sensory-motor integration, suggests that biology might offer algorithms, strategies or "lessons learned" in developing applications that have to solve similar sets of problems. To the extent that biologists may not yet completely understand natural systems, they might benefit from the advances in algorithmic formulations.

In order to maintain the focus on really hard problems, it is important to consider a variety of difficult applications and issues (e.g. Appendix 1). The examples include auditory and visual integration, sensory-motor integration, early detection of machine failure, medical diagnosis, etc. One specific example involves effective integration of diverse information – a key component of a successful implementation of networked web of intelligent satellites performing dynamic and comprehensive onboard integration of Earth observing sensors. How organisms use fusion and fusion-like processes is understood in some biological systems, to the point of generating principles and algorithms. Can these biologically-derived algorithms have an impact on the work of earth and computer scientists? In particular, how can we analyze the terabytes of information being gathered? Can we enable sensor webs without the use of large computers?

The program committee, led by Misha Pavel, Ph.D., of Oregon Graduate Institute of Science & Technology, assembled a multidisciplinary group of 29 individuals (Appendix 3) including investigators in engineering, theory of pattern recognition, and neuroscience. This report summarizes the meeting's discussions and conclusions and offers specific input for a potential Request for Proposals (RFP) in this area.

WORKSHOP PURPOSES:

> **To Identify Problems that require or might benefit from information fusion**
> **To Identify open basic and applied research questions**
> **To Identify multidisciplinary research directions and areas of novel solutions**

FORMAT

The multidisciplinary nature of the purpose of the workshop as well as the diversity of participants required a novel approach to meeting organization. To maximize the interchange of information, the meeting's agenda was left relatively open and the participants and organizers modified it as needs required. The agenda consisted of several talk sessions interspersed with breakouts (see Appendix 4 for abstracts). During the first day, survey and data talks generated a list of ideas and topics that were posted around the room and gathered by the chair. During agenda setting on day two, participants were randomly assigned to five groups. The groups generated questions that they later discussed in the meeting; participants could join the group discussing the questions they found most interesting. In formulating a written record of their discussions, participants were asked to frame questions so as to form a basis for a Request for Proposals (RFP). See Schedule in Appendix 2. This format resulted in a long and fertile list of issues. Specific issues are summarized in Appendix 1 and captured in the discussion below.

An important component of the meeting's agenda was the formation of breakout groups to discuss and identify questions, which might form the basis for an RFP. These questions and a summary of discussion of each follows.

SUMMARY OF RECOMMENDATIONS

1. **To encourage cross-fertilization** and take advantage of developments in the understanding of biological data fusion, it is important to identify tasks at a micro level, i.e., subtasks that may be part of a larger task that do have commonality among disciplines. A challenging interdisciplinary research area is to identify such basic subtasks and investigate how successes in one discipline can be applied to the other. Defining a broad range of subtasks where experience from biological data fusion can be applied to other disciplines at all levels of processing will be useful.

2. **There is a need for approaches that reduce the data collecting and archiving** requirements while increasing the signal in the data being collected and there is a need to further enhance the capability of collecting more relevant data describing

interesting events and omitting the details of non-interesting and non-events.

3. **The special nature of sensor fusion problems calls for special solutions.** Structural commonalities between sensor fusion problems in earth sciences (e.g. estimating landscape variables influencing ecosystem production models), intelligent signal processing applications (e.g. speech/vision), and basic problems of neuroscience (e.g. understanding multi-modal receptive field properties), suggest that approaches developed in one domain could provide insights into problems in other domains. Major new research questions in sensory fusion include (1) optimal sensor design with redundancy and complementarity to maximize information utility and robustness, (2) adaptive mechanisms for online selection, rejection, or modulation of sensor properties based on estimates of reliability/utility, (3) methods to exploit domain-dependent forward models to be utilized in variable estimation, (4) application of existing techniques, such as "mixture of experts", to identify and exploit domains of relevance of different sensor planes in large scale sensor fusion problems, (5) development of new principles of dynamical system design for iterative fusion/computation in multi-modal data sets.

4. **Sequential Fusion**: uncertainty driven data selection and integration. This element is concerned with systems that have some level of control over their sensory input, and can make decisions as to which data to collect and integrate. There is a formal theoretical framework, including work in Operations Research, that addresses the general problem in decision-theoretic terms, but the specific problems of large data sets, especially when the connection between the observations and the ecological parameters of interest is ill defined, remains to be addressed. The representation of the remaining uncertainty in the estimate of the quantity of interest is a key element in this decision making.

5. **Development of Metrics to assess quality of Information Fusion:** To determine the utility of data fusion, it is important to have objective measures of the contributions to important information by data fusion techniques or by data provided by additional sensors. Thus performance metrics are a function of two variables – data fusion techniques and sensor measurements. Metrics need not necessarily be absolute measurements but be relative measures of improvement. The metrics can be generic (i.e., applicable to a broad class of tasks) or task-dependent. Both types of metrics are useful and some examples of them are given below. **Generic Metrics** include: **Correlation of additional dimension of information** with other previously obtained information; **Information theory based data compression** (measure of redundancy); **Information channel** dependent uncertainty (important to know the quality of the additional channel data (its level of uncertainty) when added as an additional sensor input (will it add value?); and **Data consistency**, e.g. how different are different sensors' inputs when they measure same parameters? **Task-specific metrics** include accuracy of result (e.g., variance of estimates, classification error rates) and accomplishment of task objectives

4

6. **Issues of particular relevance to Earth Observing operations: ExploringAutomatic Registration and Georectification.** There is a need for automated registration of image and other data from multiple satellites and other sources of data. This registration involves not only accurate location of each pixel on the geoid (geolocation), but also the projection of the image into two dimensions with minimal distortion given an end-user's problem. Currently the system is largely automated, but 1) it often requires a manual step to achieve the maximum precision and 2) errors are added anytime re-projection is required. With the expected availability of vast additional data from new sensors in the future, the complete automation of this process is necessary both for speed and practicality reasons. The data being registered may include satellite imagery in multiple spectral bands at multiple spatial scales, as well as data from other sensors including air- and ground-based collection. Characterizing uncertainty in fused imagery due to integrating data across spatial resolutions

DISCUSSION POINTS: ELEMENTS AND RECOMMENDATIONS

What is common to all data fusion tasks? What is unique to specific tasks? What kinds of non-biological tasks benefit from biology? At what level must inputs be fused?

At the "macro" level, data fusion tasks are not common among disciplines. For example, climate modeling, earthquake analysis, farm productivity predictions, ecological forecasting etc. are data fusion activities in Earth sciences that do not necessarily have analogs in biology. **To encourage cross-fertilization and take advantage of developments in the understanding of biological data fusion, it is important to identify tasks at a micro level, i.e., subtasks that may be part of a larger task that do have commonality among disciplines.** A challenging interdisciplinary research area is to identify such basic subtasks and investigate how successes in one discipline can be applied to the other.

Some example subtasks that may be common across disciplines (especially between biological data fusion and others) are given below:
- Extraction of required/specific information (for a given task)
- Feature identification/classification
- Anomaly detection
- Quantitative estimation
- Adaptive learning
    - rapid redirection of attention/focus
    - adding or subtracting information channels
    - feedback of results to fusion techniques to improve performance

Focus is on the automated analysis of multiple input channels in contrast to data fusion for visualization, which will not, in general, be biologically inspired. Both biological and Earth science fusion demonstrate a progression of fusion products (small sets of sensors

5

produce intermediary inputs to higher level fusion processes). Thus, **defining a broad range of subtasks where experience from biological data fusion can be applied to other disciplines at all levels of processing will be useful.**

How would bio-inspired adaptive multi-sensor fusion provide added utility to Earth observing systems?

There is a class of problems that arise when observing complex, heterogeneous systems which can be summarized as having both too much data and not enough data. The ability to observe and characterize a complex system requires sensing technologies that often produce large volumes of data beyond our current ability to understand the data in near real time or cost effectively. There is a bottle neck in sifting and extracting useful information in detail sufficient to character the system dynamics. Reducing the data volume by degrading the sensing capability can often omit important information while maintaining high fidelity observational capabilities overwhelms our ability to off load and understand the data. Thus, **there is a need for approaches that reduce the data collecting and archiving requirements while increasing the signal in the data being collected and there is a need to further enhance the capability of collecting more relevant data describing interesting events and omitting the details of non-interesting and non-events.**

Biological systems exhibit behavior and capabilities that seem to cope with the monitoring and focusing of attention on interesting events. Biological systems are evolved to minimize uncertainty in the environment as well as respond in real time to changing environments. Possible solutions to the signal-to-noise and data volume problems would include principles of sensor integration derived from biological organisms that would maximize signal detection while minimizing false alarms. Implementations might include smart sensors, adaptable sensors, near real-time response, optimizing the logic of sensors sequence and type for cross mission tip-off and cueing (orbit, sensor resolution, and etc). There is interest in methods that can automatically focus data collection on anomalies from either a long term baseline or short term baseline. The types and resolutions of sensors, the order they are deployed, logic to detect coincidence, adaptiveness to short and long time periods are important characteristics of an implementation as is the distribution of intelligence (computation) in the system.

Sensor Fusion

Real world data analysis problems involve mappings from measured data variables, d_i, to interpretation or goal vector I. Sensor fusion problems are a specialization of this general class of mapping problems in which (1) the original data variables are grouped into high-dimensional "channels" D_i often in the form of spatial maps (e.g. pixel maps), (2) the channels can be rendered "compatible" with each other through an initial transformation/registration process (as in normalization of images through shift, rotation, and scaling) which brings the channels into approximate component by component isomorphism, (3) corresponding components (e.g. pixels) of the registered

sensory planes have complex statistical dependencies often with complementary domains of reliability (e.g. a visual cue more reliable in daytime vs. auditory cue more reliable in dark), and (4) the goal variables are nontrivial (usually nonlinear) functions of the registered sensory components, and (5) the goal variables can be used (iteratively) to improve the registration/normalization step. Output goals might include classification, regression, detection or compression which would follow the respective normalization or transformation operations.

Canonical sensor fusion problems include (1) construction of high quality composite visual images combining IR and radar signals, (2) extraction of ecosystem variables such as leaf area index (LAI) combining optical and microwave regime satellite and airborne imagery, (3) boosting spatial-temporal resolution of brain imaging combining fMRI, EEG, and MEG, (4) construction of a target location map combining visual, auditory, and somatosensory cues, and (5) extraction of shape contours in natural images from intensity, color, texture, motion, and depth cues.

The special nature of sensor fusion problems calls for special solutions. Structural commonalities between sensor fusion problems in earth sciences (e.g. estimating landscape variables influencing ecosystem production models), intelligent signal processing applications (e.g. speech/vision), and basic problems of neuroscience (e.g. understanding multi-modal receptive field properties), suggest that approaches developed in one domain could provide insights into problems in other domains. Major new research questions in sensory fusion include (1) optimal sensor design with redundancy and complementarity to maximize information utility and robustness, (2) adaptive mechanisms for online selection, rejection, or modulation of sensor properties based on estimates of reliability/utility, (3) methods to exploit domain-dependent forward models to be utilized in variable estimation, (4) application of existing techniques, such as "mixture of experts", to identify and exploit domains of relevance of different sensor planes in large scale sensor fusion problems, (5) development of new principles of dynamical system design for iterative fusion/computation in multi-modal data sets.

Sequential Fusion: uncertainty driven data selection and integration

This element is concerned with systems that have some level of control over their sensory input, and can make decisions as to which data to collect and integrate. For example, detection of forest fires can be performed on the basis of observing elevated temperatures in a region, but temperature measurement over large areas is at a very coarse resolution. How can a system best decide from the limited information in a low spatial-resolution temperature scene, which additional measurements will allow it to efficiently decide whether the small anomaly observed is indeed a forest fire, or when it can safely conclude that no further measurements are needed? There is a formal theoretical framework, including work in Operations Research, that addresses the general problem in decision-theoretic terms, but the specific problems of large data sets, especially when the connection between the observations and the ecological parameters of interest is ill

defined, remains to be addressed. The representation of the remaining uncertainty in the estimate of the quantity of interest is a key element in this decision making.

What performance metrics can be defined for data fusion?
- **Metrics of data redundancy**
- **Cost/benefit analysis**

To determine the utility of data fusion, it is important to have objective measures of the contributions to important information by data fusion techniques or by data provided by additional sensors. Thus performance metrics are a function of two variables – data fusion techniques and sensor measurements. Metrics need not necessarily be absolute measurements but be relative measures of improvement. The metrics can be generic (i.e., applicable to a broad class of tasks) or task-dependent. Both types of metrics are useful and some examples of them are given below.

Generic Metrics
- Correlation of additional dimension of information with other previously obtained information
- Information theory based data compression (measure of redundancy)
- Information channel dependent uncertainty (important to know the quality of the additional channel data (its level of uncertainty) when added as an additional sensor input (will it add value?)
- Data consistency (data resolution, temporal issues) – how different are different sensors' inputs when they measure same parameters?

Task-Specific Metrics
- Accuracy of result (e.g., variance of estimates, classification error rates)
- Accomplishment of task objectives

Issues of particular relevance to work in Earth Observing Systems:

Exploring Automatic Registration and Georectification

There is a need for automated registration of image and other data from multiple satellites and other sources of data. This registration involves not only accurate location of each pixel on the geoid (geolocation), but also the projection of the image into two dimensions with minimal distortion given an end-user's problem. Currently the system is largely automated, but 1) it often requires a manual step to achieve the maximum precision and 2) errors are added anytime re-projection is required. With the expected availability of vast additional data from new sensors in the future, the complete automation of this process is necessary both for speed and practicality reasons. The data being registered may include satellite imagery in multiple spectral bands at multiple spatial scales, as well as data from other sensors including air- and ground-based collection.

Specify a target precision relative to current standards for future Earth-observing sensors and sensor webs?

Characterizing uncertainty in fused imagery due to integrating data across spatial resolutions

Satellite imagery is collected at multiple spatial resolutions. The problem of reconciling data on different spatial scales is complicated by the fact that different statistics are expected on different scales. Because the uncertainty about each measurement depends on the spatial scale, merging measurements at different scales requires a careful treatment of uncertainty. In order to solve this problem, it is necessary first to characterize the uncertainty that we have about the raw sensor data, and then to propagate this uncertainty through a fusion step to represent uncertainty about the final information.

Special requirements for on-demand event monitoring

Current satellite imagery is collected continuously. However, future assets may collect large amounts of high-resolution data only on request. The spectral band of data collection may be selectable. A scheme is necessary to trigger the collection of specific data on demand for special requirements. This system would make it possible to make maximal use of a network of data collection assets to investigate a particular problem of interest.

CONCLUSIONS

The general conclusion from this workshop is that the dialogue between engineering, computer science and biologist can be very useful in making important advances in all three of these areas. In particular, the engineers and computer scientists can learn from the many diverse systems that evolved in natural biological organisms. The analyses of these processes provide guidance for new directions in the study of fusion algorithms as well as for the development of engineering systems.

In a similar vein, the biological and behavioral scientist can receive guidance by quantitatively modeling the empirically observed fusion processes. The formalization and quantification of the natural fusion processes are likely to lead to new questions and new approaches to study the natural processes.

Having emphasized the commonalities among these diverse disciplines, it is important to note that each type of problem is likely to give rise to very domain-specific issues requiring unique solutions. The scientists, designers and developers must therefore consider carefully the critical aspects that are common and those that are specific to the domain at hand.

Appendices

Appendix 1: ISSUES CONSIDERED

During the workshop the participants considered many issues. The following is a partial list of the topics that arose during the presentations and discussions.

1. Objective function for fusion
2. Registration and synchrony – alignment with respect to space and time (support), e.g., conformal mapping
3. Correspondence of data from different sources – information pertaining to the same or different objects
4. Calibration – Local mapping of measured quantities (e.g. IR-TV local polarity reversals)
5. Assessment of reliability of different sources
6. Using directed graphs to analyze fusion
7. Biological information should be integrated with care (Bio-superstition)
8. Multiresolution representation - variable support
9. Assessment of relevance (validity) of different sources
10. Are there principles to derive from the biological sophisticated sensory-motor integration capability?
11. Linear vs nonlinear algorithms for fusion
12. Complexity of algorithms and speed of response
13. Independence of source and processing algorithms
14. Resource allocation
15. Multilevel of fusion
16. Use humans to fuse instead of computational algorithms
17. Map uncertain information by inaccurate sensory representation.
18. Registration and synchrony – Space and time (indexes), e.g., conformal mapping
19. Generative approach vs decision boundary
20. Task-dependency
21. Fusion for human consumption vs for machine recognition
22. Evaluation is difficult because of many specific choices to be made to generate stimuli and tasks
23. Decision as fusion
24. Number of levels of inference

Appendix 2: SCHEDULE for the WORKSHOP

April 22, 2002

9:00-10:15   Introductions and charge to group
             Misha Pavel, OGI, Steve Zornetzer, NASA Ames, Joel Davis, ONR, Joe
             Coughlan, NASA Ames
10:15-12:30 Presentations: Surveys and data talks
             Misha Pavel
             Ramakrishna Nemani
                 – *Use of Satellite Remote Sensing Data in Global Change Research*
             Frank Werblin
                 - *Biologically-Inspired Techniques for the Fusion of Images from
                 Multiple Sensors*
12:30-1:30   Lunch

1:30-5:30   Data Presentations
             Peter Cheeseman
                 - *Bayesian Inference: A General Approach to Modelling from
                 Sensor Data*
             Mark Willis and Charles Higgins
                 - *Multi-Sensory Integration in Insect Flight Navigation* and
                 Information Fusion in the Dipteran Flight Navigation System
             Tim Shaw
                 - *Human-Centered Remote Sensor Information Fusion*
             Toni Jebara
                 - *Discriminative and Generative Learning in Perception and
                 Interaction Modeling*
             William Krebs
                 - *Using an Image Discrimination Model to Predict the Detectability
                 of Targets in Color Scenes*
             Robin Morris
                 - *Combining Data with Uncertain Relationships*
             Karen Moe
                 - *Data Uncertainty Challenges in NASA's Earth Science Systems
                 and Sensor Webs*
             Sarah Graves
                 - *Data Fusion in a Data Mining Framework*

6:00-8:00   Dinner

11

Tuesday, April 23, 2002

9:00-12:00   Agenda setting:
                 Division into groups to identify questions which might form basis for RFP
                 Reconvening to discuss questions
                 Reassembling in working groups

12:00-1:00   Lunch

1:30-2:30   Progress reports

2:30-5:30   Data Presentations
                 Barry Stein
                      - *Brain Mechanisms for Integrating Information from Different Senses*
                 Thomas Anastasio
                      - *A Computational Model of the Development of the Cortico-Tectal Pathways Mediating Multisensory Enhancement*
                 Stéphane Viollet
                      - *Visual/Inertial Sensory Fusion on Board a Micro Air Vehicle*
                 Jennifer Dungan
                      - *Data Fusion in Geographic Information Science: The Relevance of Spatial Support*
                 James Houk
                      - *The Brain's "Agents" for Problem Solving*
                 Hynek Hermansky
                      - *Multi-Stream Automatic Recognition of Speech*
                 Jeff Bilmes
                      - *Statistical Modeling of Data-Fusion for Classifier Systems*

5:30-6:00   Convening with group leaders

Wednesday, April 24, 2002

8:30-11:00   Compilation of reports/distribution

                 Data Presentations:
                 Rufin VanRullen
                      - *Preattentive Visual Processing of Complex Natural Scenes*
                 Nikolai Shabonov
                      - *Fusion and Uncertainties of Angular and Spectral Information in Remote Sensing of Land*
                 Laurence Maloney
                      - *Cue Combination in Biological Vision*

11:00-12:00   Closing discussion:
                 Review of Compilations
                 Next Steps

Appendix 3: Participants

Anastasio, Thomas J.
University of Illinois at Urbana-Champaign
Beckman Institute
405 N. Mathews Avenue
Urbana, IL 61801
tel: 217.244.2895
fax: 217.244.5180
email: tja@uiuc.edu

Barto, Andrew G.
University of Massachusetts, Amherst
Department of Computer Science
140 Governor's Drive
Amherst, MA 01003-4610
tel: 413.545.2109
fax: 413.545.1209
email: barto@cs.umass.edu

Berger, Theodore W.
University of Southern California
Department of Biomedical Engineering
Mail Code 1451
Los Angeles, CA 90089-1451
tel: 213.740.8017
fax: 213.740.0343
email: berger@bmsrs.usc.edu

Bilmes, Jeffrey A.
University of Washington
Department of Electrical Engineering
Box 352500
Seattle, WA 98195-2500
email: bilmes@ssli-mail.ee.Washington.edu

Cheeseman, Peter
NASA/Ames Research Center
MS 269-2
Moffett Field, CA 94035-1000
tel: 650.604.4946
fax: 650.606.3594
email: cheesem@ptolemy.arc.nasa.gov

Coughlan, Joseph C.
NASA/Ames Research Center
MS 269-3
Moffett Field, CA 94035-1000
tel: 650.604.5689
fax: 650.604.6009

13

email: jcoughlan@mail.arc.nasa.gov

Davis, Joel L.
Office of Naval Research
ONR 342
800 N. Quincy Street
Arlington, VA 22217-5660
tel: 703.696.4744
fax: 703.696.1212
email: davisjl@onr.navy.mil

Dungan, Jennifer L.
NASA/Ames Research Center
MS 242-4
Moffett Field, CA 94035-1000
tel: 650.604.3618
fax: 650.604.4680
email: jdungan@mail.arc.nasa.gov

Felzer, Benjamin
The Ecosystems Center
Marine Biological Laboratory
7 MBL Street, SB-118
Woods Hole, MA 02543
tel: 540.289.7748
fax: 508.457.1548
email: bfelzer@mbl.edu

Graves, Sara J.
University of Alabama, Huntsville
Information Technology & Systems Center
Huntsville, AL 35899
tel: 256.824.6064
fax: 256.824.5149
email: sgraves@itsc.uah.edu

Hall, David L.
Penn State University
School of Information Sciences and Technology
2E Thomas Building
University Park, PA 16802
tel: 814.865.8711
fax: 814.865.6426
email: dhall@ist.psu.edu

Hermansky, Hynek
OGI School of Science & Engineering
Department of Electrical and Computer Engineering
20000 N.W. Walker Road
Beaverton, OR 97006-8921

14

tel: 503.748.1136
fax: 503.748.1406
email: hynek@ece.ogi.edu

Higgins, Charles M.
University of Arizona
Neuromorphic Vision and Robotic Systems Lab
1230 East Speedway Boulevard
Tucson, AZ 85721-0104
tel: 520.621.6604
fax: 520.621.2478
email: higgins@ece.arizona.edu

Houk, James C.
Northwestern University
Department of Physiology
303 E. Chicago Avenue, M211
Chicago, IL 60611
tel: 312.503.8219
fax: 312.503.5101
email: houk@casbah.acns.nwu.edu

Jebara, Tony
Columbia University
Department of Computer Science
CEPSR 605, Mail Code 0401
1214 Amsterdam Avenue
New York, NY 10027
tel: 212.939.7079
fax: 212.666.0140
email: jebara@cs.columbia.edu

Jennings, Diana E. (formerly Diana Blazis)
Center for Advanced Studies in the Space Life Sciences
7 MBL Street
Woods Hole, MA 02543
tel: 508.289.7535
fax: 508.289.7951
email: djennings@mbl.edu

Kicklighter, David
The Ecosytems Center
Marine Biological Laboratory
7 MBL Street, SB-118
Woods Hole, MA 02543
tel: 508.289.7490
fax: 508.457.1548
email: dkick@mbl.edu

Krebs, William
Federal Aviation Administration

AAR-100 (Room 907A)
800 Independence Avenue, S.W.
Washington, DC 20591
tel: 202.267.8758
fax: 202.267.5797
email: william.krebs@faa.gov

MacElroy, Robert D.
NASA Ames Research Center
MS 19-20
Moffett Field, CA 94035
tel: 650.604.5573
fax: 650.604.1967
e-mail: rmacelroy@mail.arc.nasa.gov

Maloney, Laurence T.
New York University
Department of Psychology
6 Washington Place, 8th Floor
New York, NY 10003
tel: 212.998.7851
fax: 212.995.4349
email: ltm@cns.nyu.edu

Mel, Bartlett W.
University of Southern California
Department of Biomedical Engineering
Mail Code 1451
Olin Hall 500
Los Angeles, CA 90089-1451
tel: 213.740.0334
fax: 213.740.0343
email: mel@usc.edu

Moe, Karen L.
NASA/Goddard Space Flight Center
Earth Science Technology Office
MS 711
Greenbelt, MD 20771
tel: 301.286.2978
fax: 301.286.2756
email: karen.moe@gsfc.nasa.gov

Moghaddam, Mahta
California Institute of Technology
Jet Propulsion Laboratory
MS 300-227
4800 Oak Grove Drive
Pasadena, CA 91109
tel: 818.354.1591
fax: 818.393.5285

email: mahta.moghaddam@jpl.nasa.gov

Morris, Robin
NASA/Ames Research Center
MS 269-2
Moffett Field, CA 94035-1000
tel: 650 604.0158
fax: 650.604.3594
email: rdm@email.arc.nasa.gov

Nemani, Ramakrishna
University of Montana
School of Forestry
Science Complex 435
Missoula, MT 59812
tel: 406.243.4632
fax: 406.243.4845
email: nemani@ntsg.umt.edu

Pavel, Misha
Oregon Health & Science University
Department of Electrical and Computer Engineering
20000 N.W. Walker Road
Beaverton, OR 97006
tel: 503.748.1155
fax: 503.748.1406
email: pavel@ece.ogi.edu

Ramapriyan, H. K. (Rama)
NASA/Goddard Space Flight Center
MS 423
Greenbelt, MD 20771
tel: 301.614.5356
fax: 301.614.5267
email: ramapriyan@gsfc.nasa.gov

Rastetter, Edward
Marine Biological Laboratory
Ecosystems Center
7 MBL Street, SB-234
Woods Hole, MA 02543
tel: 508.289.7483
email: erastett@mbl.edu

Shabanov, Nikolay
Boston University
CAS Geography Department
675 Commonwealth Avenue
Boston, MA 02215
tel: 617.353.2525
email: shabanov@bu.edu

Shaw, Tim
Penn State University
Applied Research Laboratory
P.O. Box 30
State College, PA 16804
tel: 814.865.8812
fax: 814.863.8783
email: tshaw@psu.edu

Stein, Barry E.
Wake Forest University School of Medicine
Department of Neurobiology and Anatomy
Medical Center Building
Winston-Salem, NC 27157-1010
email: bstein@wfubmc.edu

VanRullen, Rufin
California Institute of Technology
Division of Biology, 139-74
Pasadena, CA 91125
tel: 626.395.2879
fax: 626.796.8876
email: rufin@klab.caltech.edu

Vaughan, William S.
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217
tel: 703.696.4505
fax: 703.696.1212
email: vaughaw@onr.navy.mil

Viollet, Stéphane
CNRS/Univ. de la Méditerrannée
Dept. of Biorobotics, UMR Mouvement et Perception
31, chemin Joseph Aiguier
13402 Marseille France
tel: 33.4.91.16.44.17
fax: 33.4.91.22.08.75
email: viollet@laps.univ-mrs.fr

Werblin, Frank
University of California, Berkeley
Department of Molecular & Cell Biology
142 Life Sciences Addition, #3200
Berkeley, CA 94720-3200
tel: 510.642.7236
fax: 801.640.3205
email: werblin@socrates.berkeley.edu

Willis, Mark A.
Case Western Reserve University
Department of Biology
10900 Euclid Avenue
Cleveland, OH 44106-7080
tel: 216.368.4358
fax: 216.368.4672
email: maw27@po.cwru.edu

Zornetzer, Steven F.
NASA Ames Research Center
Director, Information Sciences and Technology
Mail Stop I:200-6
Moffett Field, CA 94035-1000
tel: 650.604.2800
fax: 650.604.1726
e-mail: szornetzer@mail.arc.nasa.gov

Program committee: Pavel (chair), Couglin, Davis, Zornetzer

Appendix 4: ABSTRACTS

A COMPUTATIONAL MODEL OF THE DEVELOPMENT OF THE CORTICO-TECTAL PATHWAYS MEDIATING MULTISENSORY ENHANCEMENT
Thomas J. Anastastio, University of Illinois at Urbana

Keywords: multisensory integration, development, computational modeling

Neurons in the deep layers of the superior colliculus (SC) detect targets in the environment by integrating input from multiple sensory systems. Some deep SC neurons receive input of only one sensory modality (unimodal) while others receive input of multiple modalities. Multimodal deep SC neurons exhibit multisensory enhancement, in which the response to input of one modality is augmented by input of another modality. The phenomenon of enhancement is consistent with the hypothesis that deep SC neurons use sensory input to compute the probability that a target has appeared. Multisensory enhancement can be simulated using a model in which sensory inputs are random variables and target probability is computed using Bayes' rule. Informational analysis of the model indicates that input of another modality can indeed increase the amount of target information received by a multimodal neuron, but only if the input is ambiguous. Simple neural models can implement the processing required for computing Bayes' rule. Depending on input distributions, target probability can be computed exactly by single-neuron models or accurately estimated using feed-forward neural networks. The models suggest that multisensory enhancement is a natural consequence of neural computation and could arise from ascending sensory input alone. Interestingly, neurophysiological findings show that both ascending and cortical descending inputs are required for multisensory enhancement (Jiang, Wallace, Jiang, Vaughn, and Stein, J Neurophysiol 85: 506-522, 2001).

We have constructed a two-stage model of the development of multisensory enhancement in the deep SC that includes ascending and cortical descending inputs. Both stages of the model are unsupervised and rely only on local, neurobiologically plausible mechanisms of synaptic plasticity. The first stage involves an algorithm that increases information transmission from stochastic ascending inputs to deep SC neurons. Multimodal and unimodal units emerge at this stage, and their ratio increases as the ambiguity of the input increases, as expected from informational analysis of the Bayes' rule model. The second stage involves a novel algorithm based on correlation between the activities of deep SC and cortical neurons, and on anti-correlation between cortical neurons and ascending inputs, which alters the amount by which cortical descending influences presynaptically facilitate ascending synapses. The fully trained model reproduces the experimental finding that cortical inactivation drastically reduces multisensory enhancement but has little effect on the modality-specific responses of deep SC neurons.

STATISTICAL MODELING OF DATA-FUSION FOR CLASSIFIER SYSTEMS
Jeff A. Bilmes, University of Washington

There have been many studies which have both demonstrated empirically and/or shown theoretically that the fusion of multiple classifier systems can lead to large gains in overall performance relative to the individual classifiers being combined. In many of these cases, different combination "rules" are used to combine the different classifier outputs together in order to make a single joint decision. The success of some of these combination schemes, for example, depend on the error events of each classifier being mutually independent of each other. In this talk, we make the the case that a system of classifier combination can be understood in terms of an overall statistical system that models the process of information fusion in a classifier combination algorithm. In this view, a goal of classifier combination becomes that of statistical model selection. Using the language of directed graphical models (DGMs), several representations of two common combination schemes are obtained. Furthermore, by changing the representation slightly, we arrive at model combination rules some of which have shown improvements in simple classification tasks. This talk further will describe a new dynamic graphical model based toolkit, and how it can be used to model information fusion in time series and speech recognition model.

BAYESIAN INFERENCE: A GENERAL APPROACH TO MODELLING FROM SENSOR DATA
Peter Cheeseman, NASA/Ames Research Center

Although Bayesian inference was developed from first principles as the optimal method for learning models from uncertain data over 250 years ago, ad hoc statistical and nonstatistical methods for integrating uncertain data still persist. This can partly be explained because Bayesian inference is often computationally expensive. This excuse no longer holds in general because cheap high speed computers are now available, allowing previously intractible inference problems to be solved. Bayesian inference of 3-D surface models from multiple images will be used as an example, where Bayes Theorem is used as a form of inverse computer graphics to solve the computer vision problem. That is, in computer graphics the 3-D surface model is known, and images are generated from it. In inverse graphics, a set of images of the unknown surface are given, and the goal is to find the most likely surface that would have generated the images. There are many other examples of successful use of Bayesian inference methods in science, but some areas of science do not seem to have grasped the basics. I will focus on examples of earth science estimation problems, such as estimating net primary production (NPP), or soil moisture from satellite and ground observations. The Bayesian analysis shows why these are very hard problems, and why the estimates (so called "data products") are essentially meaningless. Part of the problem in earth science is that estimates are thought of and used as data, so the question of how accurate are the estimates is not seriously considered. In the Bayesian framework, the only data are the actual instrument readings (e.g. pixel intensity, thermometer reading, etc.), all the rest (e.g. leaf area, sea surface temperature, etc.) are **estimates**, and so subject to uncertainty. In many earth science this uncertainty is so high as to make the estimates essentially useless. It is only because

21

these estimates are not checked against "ground truth" that the inaccuracy of the estimates is not obvious.

DATA FUSION IN GEOGRAPHIC INFORMATION SCIENCE: THE REVELANCE OF SPATIAL SUPPORT
Jennifer Dungan, NASA/Ames Research Center

Map overlay analysis, the combination of multiple layers of geographic data, has been the driving motivation behind digital systems for GIS. What happens when the data in two different maps are from different -sized spatial units? For some kinds of analysis, this situation may have serious consequences. I will discuss the geostatistical concept of spatial support, equivalent to the the length, area or volume that a measurement or prediction represents, and why this is important to the analysis and modeling of geographic phenomena.

DATA FUSION IN A DATA MINING FRAMEWORK
Sara J. Graves, University of Alabama, Huntsville

Data fusion requires a framework that provides means and tools to enable the integration of data from diverse sources to obtain information of greater quality. Data mining provides such a framework that utilizes data from a variety of observations, models, and simulations to derive fused information. Furthermore, this framework can be utilized to investigate/validate relationships between features in model data and real-world observations. Often the problem with data fusion is the complexity of being able to analyze and combine large sets of heterogeneous and distributed data in a timely and meaningful manner. Data mining offers powerful new approaches to generate novel information by enabling the exploration of large amounts of data of diverse content and type (textual, geospatial, temporal, etc.) for the fusion of disparate data. This presentation will describe a data mining framework for exploratory data analysis and fusion, including examples from Earth science.

MULTI-STREAM AUTOMATIC RECOGNITION OF SPEECH
Hynek Hermansky, OGI School of Oregon Health and Sciences University
International Computer Science Institute, Berkeley, California

Linguistic message in speech is coded in a sequence of speech sounds. Since days of Isaac Newton , shape of speech spectrum is accepted as the prime carrier of this message.Subsequently, conventional automatic speech recognizer estimates the message from a temporal evolution of spectral envelopes of short-term speech spectrum, each envelope representing frequency response of a vocal tract producing a given speech sound. However, results of perceptual experiments show that hearing seems capable of independent processing of information in different frequency bands. Such independence could account for relative robustness of speech information transfer in noise. The talk elaborates on this notion and discusses an alternative approach to automatic recognition of speech in which the final decision about the message is derived by fusing information

from a number of independent information sub-streams, each sub-stream representing frequency-localized estimates of posterior probabilities of the underlying sounds of speech.

INFORMATION FUSION IN THE DIPTERAN FLIGHT NAVIGATION SYSTEM
Charles M. Higgins, University of Arizona
Keywords: sensorimotor integration, sensor guided motor behavior

In order to maintain stable flight, negotiate complex terrain, avoid obstacles and track flying targets, flies must fuse information not only from their visual system, but also from the haltere system, the ocelli, and a number of other sensory systems. Thousands of individual sensors of many types are integrated smoothly into a coherent and highly robust suite of behaviours. My talk will first address methods of fusion of information of different types within the fly's impressive visual system, and second try to expand these ideas to the fusion of visual information with other sources of information for flight navigation. I will also discuss an engineering implementation of some of these ideas to create a biologically-inspired airborne visual navigation system.

THE BRAIN'S "AGENTS" FOR PROBLEM SOLVING
James C. Houk, Northwestern University

In 1988, Marvin Minsky published The Society of Mind, in which he developed a theory of intelligence and problem solving, starting from the concept that the mind is modular. Each module functioned as an "unintelligent" agent, but, when the agents networked to form an agency, intelligence began to emerge. I will briefly review brain imaging results showing how a network of cortical areas is activated in combination to solve difficult problems. Then I will summarize new neuroanatomical findings showing that most areas of the frontal cortex have prominent loops through the basal ganglia and cerebellum, in addition to the cortical-cortical connections that are frequently contemplated. Next, I will review the neurophysiology of these subcortical loops and analyze their signal processing operations. Finally, I will propose that the neural substrate of Minsky's agent is a distributed processing module (DPM) formed by one cortical area together with it's loops through basal ganglia and cerebellum. By considering an agency formed by a collection of DPMs, one can begin to analyze the neural substrate of intelligence. A computational architecture based on the proposed theory may be well suited for combating uncertainty with fusion.

DISCRIMINATIVE AND GENERATIVE LEARNING IN PERCEPTION AND INTERACTION MODELING
Tony Jebara, Columbia University, Tommi S. Jaakkola, MIT

Keywords: Discriminative Learning, Generative Models, Time Series Prediction, Computer Vision, Audition, Perception, Imitative Behavior Learning, Feature Selection, Support Vector Machines, Maximum Entropy, Bayesian Networks

Many machine learning systems can be cast into two key paradigms: generative and discriminative learning. A generative probabilistic distribution is a principled way to model many learning and perception problems. Therein, one provides domain specific knowledge in terms of structure and parameter priors over the joint space of all variables in a system. Bayesian networks and Bayesian statistics provide a rich and flexible language for specifying this knowledge and subsequently refining it with data and observations. The final result is a distribution that is a good generator of novel exemplars or samples from the system. We present numerous examples of computer vision and perception systems that are fundamentally generative. Conversely, discriminative techniques such as support vector machines adjust a possibly non-distributional model to data optimizing for the specific task at hand (i.e. classification or prediction). This typically leads to superior performance yet compromises the flexibility of generative modeling.

Maximum Entropy Discrimination (MED) is a novel framework that elegantly combines the complementary strengths of both discriminative estimation and generative probability densities. Calculations involve distributions over parameters, margins, and priors and are provably solvable for the exponential family. Extensions include regression, feature selection, and transduction. SVMs are also naturally subsumed and can be augmented with, for example, feature selection, to obtain substantial improvements. To extend to mixtures of exponential families, we also propose a discriminative variant of the Expectation-Maximization (EM) algorithm for latent discriminative learning using a novel reverse-Jensen inequality. These two new tools extend discriminative learning to the wide range of contemporary generative models.

Empirical results on standardized data sets demonstrate the viability of the proposed hybrid discriminative-generative approach over state of the art discriminative or generative approaches in isolation. These results include domains with many irrelevant features or sources of information and thus a discriminative solution is necessary to capture and combine the critical aspects of the data. However, a generative formalism is equally crucial to specify incomplete models, prior assumptions and facilitate estimation. Subsequently, we discuss the application of discriminative-generative learning on a behavior prediction task where two audio-visual time series of humans interacting are used to implement imitative learning. Discriminative prediction helps focus resources on the task at hand, synthesizing time series predictions as a virtual animated character despite the differences in scale, frequency and structure between audio and visual signals.

USING AN IMAGE DISCRIMINATION MODEL TO PREDICT THE DETECTABILITY OF TARGETS IN COLOR SCENES

William K. Krebs, Federal Aviation Administration
Naval Reserves, Naval Research Laboratory Science and Technology 206
Albert J. Ahumada, Jr. NASA Ames Research Center

Sensor fusion combines images from multiple sensors into a single display, with the aim of enhancing operators' target detection and situational awareness in high-workload environments. Numerous researchers (Toet & Walraven, 1996; Therrien, Scrofani, & Krebs, 1997; Waxman et al., 1997; Scribner, et al., 1998) have proposed sensor-fused algorithms that will perform equal to or better than single-band imagery, however human performance studies have shown mixed results of the benefits of fusion over single-band imagery (Krebs & Sinai, in press). Objective: The goal of our research is to develop an image discrimination model that can predict the detectability of targets in color scenes. This model would be a fairly robust and good predictor of the detectability of a target in a sensor-fused scene without the need of performing a human performance study. Methods: Images were collected with visible, mid-wave, and long-wave infrared sensors and then combined by an image fusion algorithm. Observers' reaction times and accuracy scores were collected in a variety of visual search tasks using single and dual-band imagery as well as an image discrimination model was developed to predict the effects of masking of luminance and chromatic targets by color variations in the background scene. Results: Visual search results found that sensor fusion did not improve performance relative to that obtained with single-band imagery on a target detection task. Moreover, these experiments demonstrate significant masking of color targets by color variations in the background texture. Conclusions: Actual or potential applications of this research include a quantitative methodology to evaluate the performance of an image-fused algorithm for automobile, aviation, and maritime applications.

CUE COMBINATION IN BIOLOGICAL VISION
Laurence T. Maloney, New York University

Keywords: cue combination, depth perception, shape perception, color perception, statistical decision theory, Bayesian vision

There are a variety of problems in biological vision that are commonly modeled as cue combination ('fusion') problems, the most obvious being that of depth perception. Each cue is the result of a modular computation applied to retinal information and the fusion problem involves combining the different cues to produce an overall representation of, for example, depth. I will discuss a small number of issues concerning biological cue combination in the familiar case of depth perception and in the less familiar case of surface color perception and describes experiments by myself and other intended to characterize biological cue combination in human vision. In particular I will describe an ongoing set of experiments (with Ipek Oruc) that tests whether human observers respond to changes in cue reliability optimally.

Our results together with previous work suggests that the human brain is a half-way decent statistician and that optimal models of cue combination drawn from the fusion literature are worth considering as hypotheses concerning human visual functioning.

CUE FUSION HELPS AUTOMATIC CONTOUR/SURFACE EXTRACTION
Bartlett W. Mel. University of Southern California

A key function of the primate visual system is to recognize objects and understand scenes. Empifically, a cartoon-like representation of an image, consisting of an outline drawing with simple color fills, minimizes information content while maximizing image intelligibility for object or scene classification. Automatic extraction of shape-defining contours and surface properties are very difficult tasks, however, in that (1) both contour and surface computations depend on several types of visual cues acting over long distances in the image, and (2) the rules for optimally combining evidence from many varied sources are complex and highly nonlinear. The resulting computation can be expressed as a nonlinear dynamical system operating in very high dimension. We have developed an approach to the dual contour/surface extraction problem, which uses (1) strong representational biases in the network architecture involving probabilistically-derived knowledge of the main cue interaction nonlinearities, and (2) a learning rule which sets the remaining unknown network parameters based on specially designed training sets. We report significant progress towards the automatic construction of cartoon-like representations, and demonstrate the performance of our network applied to complex visual scenes. We discuss the implications for long-range contextual processing in primate visual cortex.

DATA UNCERTAINTY CHALLENGES IN NASA'S EARTH SCIENCE SYSTEMS AND SENSOR WEBS

Karen L. Moe, NASA/Goddard Space Flight Center

Remote sensing systems produce huge volumes of data and information daily today. Within the first year of the Terra spacecraft launch in Dec. 2000, NASA's entire Earth science holdings, formerly 284 terabytes, doubled. NASA's vision for Earth science pushes the acquisition and delivery of new data and information via an intelligent web of space-based sensors, processing networks, and distribution systems. A major challenge is that greater understanding of the scientific basis for Earth system behavior and response is required in order to intelligently acquire and interpret data from sensor webs. Revolutionary advances in both science and technology are essential to achieving this vision.

This presentation will highlight some of the current challenges and future concepts that will drive the need for advances in data fusion and related enabling technologies. New observation strategies are contemplated to support remote sensing from a variety of vantage points, such as low Earth or geostationary orbiting spacecraft, sentinel sensors at libration points, in situ sensors and unmanned vehicles in the atmosphere and oceans. Both interactive and autonomous methods are needed to interpret the situation and proactively respond. This response is of two types and purposes. One is to determine the next course of action so that appropriate and full observations are made, and the other is to provide processed data and information products to scientists and the user community

in near real time. These concepts suggest major breakthroughs are needed, especially in the ways that data and information are organized for discovery and transformation.

The presentation will also suggest some of the questions and steps that the computer and information sciences community should address in order to achieve this vision. The goal is to stimulate dialogue on the similarities and differences between NASA's Earth science perspective with other data fusion applications and explore what technology thrusts might be waiting to be exploited or influenced to help address these challenges.

DATA FUSION ACROSS ELETROMAGNETIC SPECTRUM TO REDUCE UNCERTAINTY IN TARGET VARIABLE ESTIMATION
Mahta Moghaddam, Jet Propulsion Laboratory, California Institute of Technology

Electromagnetic (EM) waves with wavelengths ranging from microns to decimeters are widely used from a multitude of spaceborne, airborne, and ground-based remote sensors to generate information about a variety of targets and their dynamics. The targets include vegetated ecosystems, arid/semiarid lands, wetlands, oceans, atmosphere, urban areas, and other surface and subsurface man-made structures. Each of these target types is typically described by a large number of variables that characterize its location, geometry, and the material makeup. Remote active and/or passive measurements using EM waves enable, in principle, estimation of one or more of the pieces of information needed to fully characterize the targets in a nondestructive fashion. EM waves of various frequencies respond differently to various target characteristics. Hence, ideally, remotely sensed EM data over a large range of frequencies is required to adequately characterize the required target variables over a useful range of values. From a mathematical point of view, combining frequency-diverse EM data expands the dimensionality of the observation space from which a large number of variables can be estimated. If each data dimension (frequency) is sufficiently sensitive to one or more of the unknown variables through a known relationship, the variables can be estimated with reduced uncertainty as a larger number of observations is included. This presentation describes the nonlinear estimation of an arbitrary number of unknowns from fusion of several remote sensing data types, with a specific example for a vegetated ecosystem. A similar methodology can be applied to other disciplines to derive appropriate system attributes from multiple independent measurements of a given system.

COMBINING DATA WITH UNCERTAIN RELATIONSHIPS
Robin D. Morris and Vadim N. Smelyanskiy , NASA/Ames Research Center

The consistent treatment of uncertainty is fundamental to the correct fusion of different data streams -- without knowing the relative weightings that we need to give to each data stream, we cannot know how to correctly cmbine them. Uncertainty is represented by probability distributions over the parameters of interest, and the laws of probability theory, especially Bayes' theorem, provide the tools for combining these distributions.

However it is not necessarily straightforward to determine the forms of the distributions, and computing with these distributions may require techniques that are unfamiliar to those working in the application domain, or even the development of new computational approaches.

I will present an example from the Earth Science domain, using two data sets. The first is information from a survey-based soil database and the second are satellite observations. The information in the soil database is used to construct prior distributions over the Available Water Capacity (AWC) in the soil. Satellite observations can be used to estimate Leaf Area Index (LAI), and LAI can be related to AWC. Using these relations, the prior distribution over soil AWC can be updated to take into account the information in the satellite data. This update must take into account all sources of uncertainty, and because of the large uncertainty in the relationships between the satellite observations and LAI, and LAI and AWC, the reduction in uncertainty achievable in the soil AWC is limited.

USE OF SATELLITE REMOTE SENSING DATA IN GLOBAL CHANGE RESEARCH
Ramakrishna Nemani, University of Montana

Keywords: Global change, Earth Observing System, Earth system science

Data collected by orbiting satellites has greatly benefited global change research in recent decades. Helping this research, more than a dozen sensors are currently in operation and more are on the way. While data from individual sensors contributed to important discoveries, as in the case of Advanced Very High Resolution Radiometer for detecting changes in global vegetation, the synergistic use of data from various orbiting sensors is in its infancy. The Earth Observing System, consisting of multiple sensors on a single platform, is conceived with the explicit objective of observing various components (land, ocean and the atmosphere) of the Earth system simultaneously. While the science behind EOS is mature, the operational aspects of data fusion from different sensors are not.

Effective integration of data from various spatial, spectral and temporal domains has been found to reduce uncertainty and improve information content as in the case of mapping global land cover. However, mapping and monitoring dynamic phenomenon such as wildfires require a level of sophistication that is lacking in the current processing systems. A number of factors currently preclude active fusion of data from various sensors. These include rigid algorithms, asynchronous processing schedules and brittle architectures. Nearly two years of data from EOS/TERRA sensors should be a catalyst for developing tools for active data fusion.

NASA plans to launch a number of smaller missions forming a sensor-web in the coming decades. To succeed such a concept requires a quantum leap in our ability to process, integrate, and understand the data so that the system can adapt in near-real time fashion.

28

FUSION AND UNCERTAINTIES OF ANGULAR AND SPECTRAL INFORMATION IN REMOTE SENSING OF LAND

Nikolay V. Shabanov, Yuri Knyazikhin and Ranga B. Myneni, Boston University

Key words: remote sensing, fusion of spectral and angular information, input data and model uncertainties

Advances in remote sensing technology and radiative transfer modeling greatly improved the possibility of accurate estimates of biophysical information from spatial, spectral, angular and temporal dimensions of remotely sensed data. The retrieval technique for these parameters is generally an ill-posed problem, which implies that specification of uncertainty of data and corresponding models is required to derive solution.

This presentation discusses the benefits of fusion of remotely sensed data from different domains and highlights limitations due to accumulation of associated uncertainties on the basis of two case studies. In first case study, combination of angular and spectral information (angular signatures in spectral space) is introduced to characterize different land cover types. New metrics include angular signature slope, length and intercept. The statistical analysis with these indices confers the idea that incorporation of the directional variable should improve traditional land cover classification based on spectral information only. The second case study asses the of influence of uncertainties in spectral information on retrieval quality of Leaf Area Index (LAI). The uncertainties in the land surface reflectances and radiative transfer model are used in the algorithm to determine the quality of the retrieved LAI fields. When the amount of spectral information input to the retrieval technique is increased, not only does this increase the overall information content but also decreases the summary accuracy in the data. The former enhances quality of the retrievals, while the latter suppress it. The total uncertainty sets a limit on the quality of the retrieved fields. A stabilizing uncertainty is introduced, which is basic information to the retrieval technique required to establish its convergence; that is the more the measured information and the more accurate this information is the more reliable and accurate the algorithm output will be.

HUMAN-CENTERED REMOTE SENSOR INFORMATION FUSION

Timothy S. Shaw, Eileen S. Rotthoff, Amulya K. Garga, Penn State Applied Research Laboratory, David L. Hall, Rashaad E. Jones, Penn State School of Information Sciences and Technology

Keywords: fusion, visualization, information, remote sensing, earth science, data analysis

Abstract: In recent years, extensive research has focused on the development of techniques for multi-sensor data fusion systems in many domains. Military fusion systems process data from multiple remote sensors to develop improved estimates of the position, velocity, attributes, and identity of objects such as targets or entities of interest. The amount and complexity of the earth science data collected by NASA remote sensors underscores the need for research into strategies and techniques to facilitate its analysis. This presentation will summarize our research into human-centered remote sensor

information fusion. This work is focused on two areas for improved understanding of multi-source data; (1) information fusion techniques to combine remote sensor data with other information sources for improved contextual interpretation and understanding, and (2) data visualization methods to model, represent, and display complex multi-dimensional terrestrial and atmospheric data and processes. Information fusion techniques include the correlation and conditioning of data products, both geo-spatially and temporally, and fusion and interpretation of data using a hybrid reasoning approach. Data visualization techniques include the application of immersive, multi-modal (visual and aural) human-computer interfaces. Research for military systems will be summarized and current work with MODIS earth science products from NASA's Terra satellite will be discussed. We believe that human-centered fusion methods will provide aids to reduce cognitive biases, improve the understanding of heterogeneous, multi-source data, and provide increased opportunities for data discovery.

BRAIN MECHANISIMS FOR INTEGRATING INFORMATION FROM DIFFERENT SENSES
Barry E. Stein, Wake Forest University

Keywords: Multisensory Integration, Midbrain, Cortex, Biology

Midbrain neurons in the superior colliculus (SC) are able to synthesize information from different senses, thereby substantially altering their responses to visual, auditory and somatosensory stimuli. In some circumstances the multisensory responses are significantly enhanced so that they exceed the neuron's response to a stimulus from the most effective sense and can exceed the sum of the responses to both senses. In other circumstances they are depressed so that normally vigorous responses can be eliminated. These neural changes are paralleled by changes in SC-mediated orientation behaviors. One might reasonable hypothesize that this capability would be present in every neuron receiving convergent inputs from two or more senses. But, we have found that descending influences from cortex are essential for this process. When these corticotectal influences are removed, SC neurons lose their multisensory integration abilities, yet retain their ability to respond to cues from multiple sensory modalities. Two cortical areas have been identified as critical in this process: the anterior ectosylvian sulcus (AES) and the rostral lateral suprasylvian sulcus (rLS). These observations outline a critical brain circuit for multisensory integration that may serve as a model for the construction of silicone-based fusion devices. They may also provide the building blocks for developing network models of sensor fusion.

PREATTENTIVE VISUAL PROCESSING OF COMPLEX NATURAL SCENES
Rufin VanRullen, Fei-Fei Lli, Lavanya Reddy, Pietro Perona and Christof Koch
California Institute of Technology

Attention is often conceptualized as a reflection of the computational "bottleneck" between early and late stages of visual processing: only low-level stimulus features can be processed preattentively, leading to "pop-out" effects in visual search, while higher-level aspects (e.g. feature conjunctions, object identity) can only be registered after the serial recruitment of attentional resources. In contradiction with this view, our recent

30

results indicate that visual categorization of natural scenes (e.g. "animal vs non-animal", or "vehicle vs non-vehicle") can still be performed when attention is tied away by a concurrent demanding task ("dual-task" paradigm). Thus, attention is not a prerequisite of "high-level" vision. The classical view of attention is even more challenged when we compare the results of our dual-task paradigm with performance in visual search tasks: whether or not a target pops-out of an array of distractors (visual search performance) is independent of whether or not the target and distractors can be discriminated preattentively (dual-task performance). We show examples of discrimination tasks that can be performed without attention, but for which visual search performance is "serial" (including the preceding "animal vs non-animal" task) and other tasks that lead to "pop-out" effects in visual search, but in which the target and distractors can not be discriminated without attention. Thus, the classical dichotomies "parallel vs. serial" processing and "preattentive vs. attentive" processing are not equivalent. We suggest that preattentive tasks correspond to situations where selective neuronal populations exist for targets vs. distractors, independent of the level of processing involved, while pop-out effects (i.e. "parallel" processing) can also rely on center-surround, texture or figure-ground grouping mechanisms in early stages of the visual cortex. Neuronal receptive fields throughout the visual system can act as independent processing channels, and there is no need for attentional processes as long as (i) the target stimulus can be "picked up" by at least one selective channel and (ii) each channel contains only one stimulus.

VISUAL/INERTIAL SENSORY FUSION ON BOARD A MICRO AIR VEHICLE
Stéphane Viollet and Nicolas Franceschini, CNRS/Univ. de la Méditerrannée

Key words: Vision, flies, motion perception, optic flow, inertial sensing, micro air vehicles, tracking, sensory fusion

In the framework of our biologically inspired robotic approach, we designed, realized and tested a miniature two-propeller aerial robot which stabilizes and orients towards contrasting targets by merging visual and inertial cues. This 100 gram experimental seeing robot is tethered to a long wire secured to the ceiling of the laboratory so as to restrict its degrees of freedom to essentially motion about the yaw axis. The robot uses a novel kind of visual system, called OSCAR (Optical SCaning sensor for Autonomous Robots) 1, which was inspired by data recently obtained at our laboratory on the compound eye of flies. OSCAR associates motion detection with microscanning and has a number of outstanding features such as:

- super-resolution in target detection
- hyper-acuity in target localization
- relative invariance with contrast
- relative invariance with distance (0 - 2.5 m)

The OSCAR visual system drives the two propellers of the robot differentially so that the latter orients its gaze towards a target such as a vertical edge or bar. We added to this basic visuomotor feedback loop an inertial feedback loop based on a micro rate-gyro. The latter provides the robot with a stability augmentation system, which also improves its

31

dynamic response. In flies, the two halteres act like a gyroscope2. They drive the wing muscles and were recently shown to be under the influence of the visual motion detection system3.

Our fusion scheme for visual and inertial cues (Figure 1) is such that the error signal from the outer loop (angular orientation loop) drives the inner-loop (angular velocity loop). As a consequence, the micro-robot stabilizes in yaw with respect to its stationary optical environment. If the target happens to move, the robot keeps locking on to it and tracks it4 at angular speeds up to 27°/s – a figure close to the maximal speed for visual smooth pursuit in the human eye.

Our minimalistic robotic approach, which is largely instructed by the visuomotor and inertial control systems of flies, incites us to find solutions that make efficient use of limited material, computational and energy resources. By using purely analog neuromorphic processing we not only build small machines that keep close to biology in spirit but we also glean, in return, an interesting feedback to improve our biological understanding.

REFERENCES

[1] S. Viollet, and N. Franceschini (1999), Biologically-Inspired Visual Scanning Sensor for Stabilization and Tracking, Proceedings IEEE Int. Conf. Intelligent Robots and Systems (IROS' 99), Kyongju, Korea, pp. 204-209
[2] G. Fraenkel and J. Pringle (1938) Halteres of flies as gyroscopic organs of equilibrium, Nature 141, 919-920
[3] W. P. Chan, F. Prete and M. H. Dickinson (1998), Science, vol. 280, pp. 289-292.
[4] S. Viollet, and N. Franceschini (2001), Aerial minirobot that stabilizes and tracks with a bio-inspired visual scanning sensor, Biorobotics, B. Webb and T. Consi (Eds), MIT Press, Cambridge, pp. 67-73.

BIOLOGICALLY-INSPIRED TECHNIQUES FOR THE FUSION OF IMAGES FROM MULTIPLE SENSORS

Frank Werblin, Tibor Kozek, University of California, Berkeley

Keywords: Night vision, Image Fusion, Biological-Inspired, Image Processing

The most salient information generated by sensor arrays of different spectral sensitivity, say, IR and I2, will often lie in different spatial regions in their outputs. In order to enhance the viewer's ability to understand and interpret multi-spectral data, it is desirable to create a single consolidated image that contains all significant information obtained from the sensors. The challenge is to prioritize the information content of each sensor for each position in space and for each frame, then weight these areas of content appropriately to display the most information-rich features from both sensors in the final representation.

We have developed techniques for accomplishing this prioritization and weighting, a form of image fusion, by borrowing from techniques derived from biological image

processing. The resulting representation displays the best features of all image for each frame in a composite representation. The representation can also be color-coded to enhance scene understanding and feature detection by highlighting specific visual features such as movement or elements of the natural environment (trees, sky, etc.).

Recent studies of biological image processing reveal that the representation of illuminance for each pixel of the image is separated into values either above or below a local mean intensity. These representations are then rectified and separated into a set of space-time frequency bins. In the biological realm each space-time bin represents a distinct "feature detector." For image fusion we have borrowed the processes of above-below the mean, signal rectification and separation into spatial frequency domains. These components of the scene are then recombined, using prioritization techniques, to generate the final fused image. Many of the details of the process and the correspondence between the biology and fusion techniques will be outlined.

MULTI-SENSORY INTEGRATION IN INSECT FLIGHT NAVIGATION
Mark A. Willis, Case Western Reserve University

Animal locomotion is often studied as a model for how nervous systems generate rhythmic output and then integrate that with sensory information to shape the final expression of the behavior. Flight is one of the most challenging forms of locomotion animals have evolved. It requires all of the sensory-motor interactions of other forms of locomotion, and typically requires faster processing times to match the faster relative speeds generated in flight.

According to many rules of thumb, insects are the most successful animals to have evolved and much of their success can be attributed to their ability to fly. The insect flight system comprises functionally integrated sensory and motor systems that are specifically tuned to the dynamics of the body, and form motor-control feedback loops that are closed by the body's interaction with the environment. My laboratory studies how flying moths track wind-borne trails of odor molecules through the environment to locate distant unseen resources (e.g., food, egg-laying sites, and mates). Successful completion of this task requires that information about the odor plume modulate the processing of other sensory information and the motor output of the flight system. A prominent feature of this sensory processing (and perhaps multi-sensory processing in general) is its context dependency. Odor-tracking behavior occurs only at specific times during the night and when the animals are in specific physiological states. We have addressed this system from multiple levels of organization (i.e., free-flight behavior to neural responses) and will report on the results of experiments aimed at understanding how multimodal sensory information is used by the nervous system to generate and control adaptive behavior.